

A Psycholinguistically Motivated Parser for CCG

Michael Niv*

Technion – Israel Institute of Technology

Haifa, Israel

Internet: niv@linc.cis.upenn.edu

Abstract

Considering the speed in which humans resolve syntactic ambiguity, and the overwhelming evidence that syntactic ambiguity is resolved through selection of the analysis whose interpretation is the most ‘sensible’, one comes to the conclusion that interpretation, hence parsing take place incrementally, just about every word. Considerations of parsimony in the theory of the syntactic processor lead one to explore the simplest of parsers: one which represents only analyses as defined by the grammar and no other information.

Toward this aim of a simple, incremental parser I explore the proposal that the competence grammar is a Combinatory Categorical Grammar (CCG). I address the problem of the proliferating analyses that stem from CCG’s associativity of derivation. My solution involves maintaining only the maximally incremental analysis and, when necessary, computing the maximally right-branching analysis. I use results from the study of rewrite systems to show that this computation is efficient.

1 Introduction

The aim of this paper is to work towards a computational model of how humans syntactically process the language that they hear and read. The endpoint of this enterprise is a precise characterization of the process that humans follow, getting details such as

*The research reported here was conducted as part of my Ph.D. thesis work at the University of Pennsylvania and supported by the following grants: DARPA N00014-90-J-1863, ARO DAAL03-89-C-0031, NSF IRI 90-16592, Ben Franklin 91S.3078C-1. Preparation of this paper was supported by a postdoctoral fellowship at the Technion in Israel. I am grateful to Mark Hepple, Mitch Marcus, Mark Steedman, Val Tannen, and Henry Thompson for helpful suggestions, and to Jeff Siskind for help with typesetting CCG derivations. Any errors are my own.

timing and garden pathing exactly right.

1.1 Ambiguity Resolution

Recently, a great deal of evidence has accumulated that humans resolve syntactic ambiguity by considering the meaning of the available analyses and selecting the ‘best’ one. Various criteria for goodness of meaning have been advanced in the psycholinguistic literature: e.g. thematic compatibility and lexical selection (Trueswell and Tanenhaus 1994), discourse felicity of definite expressions (Altmann *et al.* 1994), temporal coherence in discourse (Trueswell and Tanenhaus 1991), grammatical function *vis a vis* given/new status (Niv 1993b), and general world-knowledge (Kawamoto and Farrah 1993).

Many of the works cited above consider the timing of the ambiguity resolution decision. The evidence is overwhelming that ambiguity is resolved within a word or two of the arrival of disambiguating information — that is, when there is a meaning-based criterion which militates toward one or another syntactically available analysis, that analysis is selected. Should the other analysis turn out to be the ultimately correct analysis, a garden path will result. Given that the various analyses available are *compared* on various criteria of sensibleness, it follows that these analyses are *constructed and maintained in parallel* until disambiguating information arrives. Indeed, there is psycholinguistic evidence that the processor maintains the various analyses in parallel (Nicol and Pickering 1993; MacDonland *et al.* 1992).

Our parser, therefore, must be able to build and maintain analyses in parallel. It must also extract from the developing parse in a *prompt* fashion all of the semantically relevant syntactic commitments (e.g. predicate-argument relations) in order to allow

the interpretation module that it feeds to make accurate evaluations of the meaning. Recovery from garden paths is not addressed in this paper.

1.2 Parser and Grammar

Let us adopt the widely held position that humans possess a representation of grammatical competence which is independent of any process (e.g. production, perception, acquisition) that uses it. Steedman (1994) argues that if two theories of the grammar and processor package have identical empirical coverage, but one has a more complex parser, then the other is preferred. This preference is not just on philosophical grounds of cleanliness of one's theories, but stems from consideration of the evolution of the human linguistic capacity: A theory whose grammar requires a complex parser in order to be of any use would entail a more complex or less likely evolutionary path which the parser and grammar took together than would a theory whose grammar requires little specialized apparatus by way of a parser, and could thus have evolved gradually.

So what is the simplest parser one can construct? In other words, what is the minimal addition of computational apparatus to the competence grammar necessary to make it parse? From the argument in section 1.1, this addition must include a mechanism for maintaining analyses in parallel. Minimally, nothing else is necessary — the data structure which resides in each parallel slot in the parser is a direct representation of an analysis as defined by the competence machinery.

Suppose the grammatical competence is one that always divides an English clause into a subject and a predicate (VP henceforth). Suppose also that the primary operations of the grammar are putting constituents together. Could the minimal parser for such a grammar account for the minimal pair in (1)?

- (1) a. The doctor sent for the patient arrived.
- b. The flowers sent for the patient arrived.

(1)a is a garden path. In (1)b the garden path is avoided because flowers are not good senders. The difference between (1)a and b indicates that well before the word 'arrived' is encountered, the processor has already resolved the ambiguity introduced by the word 'sent'. That is, in the main-verb analysis of 'sent', the interpreter is aware of the relation between the subject the verb *before* the end of the VP. But the minimal parser cannot put the subject together with 'sent' or 'sent for the' because the latter are not a complete VP!

There are two possible solutions to this problem, each relaxes one of the two suppositions above: Steedman (1994) argues for a grammatical theory (CCG) which does not always make the subject-

predicate juncture the primary division point of a clause. Shieber and Johnson (1993) on the other hand, argue that there is no need to assume that a constituent has to be complete before it is combined with its sister(s). At this time, neither approach is sufficiently developed to be evaluable (e.g. they both lack broad coverage grammar) so either one is viable. In this paper, I develop the first.

2 Preliminaries

CCG is a lexicalized grammar formalism — a lexicon assigns each word to one or more grammatical categories. Adjacent constituents can combine by one of a small number of combinatory rules. The universe of grammatical categories contains a collection of basic categories (e.g. atomic symbols such as *n*, *np*, *s*, etc. or Prolog terms such *np(3,sg)*) and is closed under the category-forming connectives */* and **. Intuitively a constituent of category *X/Y* (resp. *X\Y*) is something of category *X* which is missing something of category *Y* to its right (resp. left). The combinatory rules are listed¹ in table 1. They formalize this intuition. A combinatory rule may be qualified with a predicate over the variables *X*, *Y*, and *Z*₁...*Z*_{*n*}.

A *derivation* is a binary tree whose leaves are each a single-word constituent, and whose internal nodes are each a constituent which is derived from its children by an application of one of the combinatory rules. A string *w* is grammatical just in case there exists a derivation whose frontier is *w*. I equivocate between a derivation and the constituent at its root. An *analysis* of a string *w* is a sequence of derivations such that the concatenation of their frontiers is *w*.

3 The Simplest Parser

Let us consider the simplest conceivable parser. Its specification is “find all analyses of the string so far.” It has a collection of slots for maintaining one analysis each, in parallel. Each slot maintains an analysis of the string seen so far — a sequence of one or more derivations. The parser has two operations, as shown in figure 1.

¹Two common combinatory rules, type-raising and substitution are not listed here. The substitution rule (Steedman 1987) is orthogonal to the present discussion and can be added without modification. The rule for type-raising (see e.g. Dowty 1988) can cause difficulties for the parsing scheme advocated here (Hepple 1987) and is therefore assumed to apply in the lexicon. So a proper name, for example, would be have two categories: *np* and *s/(s\np)*.

- **scan**
 - get the next word from the input stream
 - for each analysis *a* in the parser’s memory
 - empty the slot containing *a*
 - for each lexical entry *e* of the word
 - make a copy *a'* of *a*
 - add the leaf derivation *e* to the right of *a'*
 - add *a'* as a new analysis
- **combine**
 - for each analysis *a* in the parser’s memory
 - if *a* contains more than one constituent
 - and some rule can combine the rightmost two constituents in *a*
 - then make a copy *a'* of *a*
 - replace the two constituents of *a'* by their combination
 - add *a'* as a new analysis

Figure 1: Parser operations

This parser succeeds in constructing the incremental analysis (2) necessary for solving the problem in (1).

$$(2) \quad \frac{\frac{\text{the}}{s/(s\backslash np)/n} \quad \frac{\text{flowers}}{n} \quad \text{sent}}{s/(s\backslash np)} \xrightarrow{>0} \frac{}{s/pp} \xrightarrow{>1}$$

But this parser is just an unconstrained shift-reduce parser that simulates non-determinism via parallelism. It suffers from a standard problem of simple bottom-up parsers: it can only know when a certain substring has a derivation, but in case a substring does not have a derivation, the parser cannot yet know whether or not a larger string containing the substring will have a derivation. This means that when faced with a string such as

- (3) The insults the new students shouted at the teacher were appalling.

the parser will note the noun-verb ambiguity of ‘insults’, but will be unable to use the information that ‘insults’ is preceded by a determiner to rule out the verb analysis in a timely fashion. It would only notice the difficulty with the verb analysis after it had come to the end of the string and failed to find a derivation for it. This delay in ruling out doomed analyses means that the parser and the interpreter are burdened with a quickly proliferating collection of irrelevant analyses.

Standard solution to this problem (e.g. Earley’s

Forward Combination				rule name	Backward Combination				rule name
X/Y	Y	\longrightarrow	X	>0	Y	X\Y	\longrightarrow	X	<0
X/Y	Y Z	\longrightarrow	X Z	>1	Y Z	X\Y	\longrightarrow	X Z	<1
X/Y	Y Z ₁ Z ₂	\longrightarrow	X Z ₁ Z ₂	>2	Y Z ₁ Z ₂	X\Y	\longrightarrow	X Z ₁ Z ₂	<2
	\vdots					\vdots			
X/Y	Y Z ₁ ... Z _n	\longrightarrow	X Z ₁ ... Z _n	>n	Y Z ₁ ... Z _n	X\Y	\longrightarrow	X Z ₁ ... Z _n	<n
	\vdots					\vdots			

|Z stands for either /Z or \Z. Underlined regions in a rule must match.

Table 1: The combinatory rules

1970 parser; LR parsing, Aho and Johnson 1974) consider global properties of the competence grammar to infer that no grammatical string will begin with a determiner followed by a verb. These solutions exact a cost in complicating the design of the parser: new data structures such as dotted rules or an LR table must be added to the parser. The parser is no longer a generic search algorithm for the competence grammar. Given the flexibility of CCG derivations, one may consider imposing a very simple constraint on the parser: every prefix of a grammatical string must have a derivation. But such a move is too heavy-handed. Indeed CCG often gives left-branching derivations, but it is not purely left-branching. For example, the derivation of a WH-dependency requires leaving the WH-filler constituent uncombined until the entire gap-containing constituent is completed, as in (4).

(4)

whose	cat	did	Fred	find
$\frac{q/(s \backslash np)/n}{q/(s \backslash np)}$	$\frac{cat}{n}$	$\frac{did}{s/s}$	$\frac{Fred}{s/(s \backslash np)}$	$\frac{find}{s \backslash np/np}$
>0		>1		
		>1		
s/np				
>0				
q				

4 The Viable Analysis Criterion

Given the desideratum to minimize the complexity of the biologically specified parser, I propose that the human parser is indeed as simple as the scan-combine algorithm presented above, and that the ability to rule out analyses such as determiner+verb is not innate, but is an acquired skill. This ‘skill’ is implemented as a criterion which an analysis must meet in order to survive. An infant starts out with this criterion completely permissive. Consequently it cannot process any utterances longer than a few words without requiring excessively many parser slots. But as the infant observes the various analyses in the parser memory and tracks their respective

outcomes, it notices that certain sequences of categories *never* lead to a grammatical overall analysis. After observing an analysis failing a certain number of times and never succeeding, the child concludes that it is not a viable analysis and learns to discard it. The more spurious analyses are discarded, the better able the child is to cope with longer strings.

The collection of analyses that are maintained by the parser is therefore filtered by two independent processes: The Viable Analysis Criterion is a purely syntactic filter which rules out analyses independently of ambiguity. The interpreter considers the semantic information of the remaining analyses in parallel and occasionally deems certain analyses more sensible than their competitors, and discards the latter.

Given that English sentences rarely require more than two or three CCG constituents at any point in their parse, and given the limited range of categories that arise in English, the problem of learning the viable analysis criterion from data promises to be comparable to other n-gram learning tasks. The empirical validation of this proposal awaits the availability of a broad coverage CCG for English, and other languages.²

5 CCG and flexible derivation

5.1 The Problem

CCG’s distinguishing characteristic is its derivational flexibility — the fact that one string is potentially assigned many truth-conditionally equivalent analyses. This feature is crucial to the present approach of incremental parsing (as well as for a range of grammatical phenomena, see e.g. Steed-

²In addition to the category-ambiguity problem in (3), the viable analysis criterion solves other problems, analogous to shift-reduce ambiguities, which are omitted here for reasons of space. The interested reader is referred to Niv (1993a) for a comprehensive discussion and an implementation of the parser proposed here.

man 1987, 1994; Dowty 1988). But the additional ambiguity, sometimes referred to as ‘spurious’, is also a source of difficulty for parsing. For example, the truth-conditionally unambiguous string ‘John was thinking that Bill had left’ has CCG derivations corresponding to each of the 132 different binary trees possible for seven leaves. The fact that this sentence makes no unusual demands on humans makes it clear that its exponentially proliferating ambiguous analyses are pruned somehow. The interpreter, which can resolve many kinds of ambiguity, cannot be used to for this task: it has no visible basis for determining, for example, that the single-constituent analysis ‘John was thinking’ somehow makes more sense (in CCG) than the two-constituent analysis ‘John’+‘was thinking’.

Note that the maximally left-branching derivation is the one which most promptly identifies syntactic relations, and is thus the preferred derivation. It is possible to extend the viable analysis criterion to encompass this consideration of efficiency as well. The infant learns that it is usually most efficient to combine whenever possible, and to discard an analysis in which a combination is possible, but not taken.³

While this left-branching criterion eliminates the inefficiency due to flexibility of derivation, it gives rise to difficulties with (5).

$$(5) \quad \frac{\text{John}}{s/vp} \quad \frac{\text{loves}}{vp/np} \quad \frac{\text{Mary}}{np} \quad \frac{\text{madly}}{vp \backslash vp}$$

In (5), it is precisely the *non*-left-branching derivation of ‘John loves Mary’ which is necessary in order to make the VP constituent available for combination with the adverb. (See Pareschi and Steedman 1987.)

5.2 Previous Approaches

Following up on the work of Lambek (1958) who proposed that the process of deriving the grammaticality of a string of categories be viewed as a proof, there have been quite a few proposals put forth for computing only normal forms of derivations or proofs (König 1989; Hepple and Morrill 1989; Hepple 1991; *inter alia*). The basic idea with all of these works is to define ‘normal forms’ — distinguished members of each equivalence class of derivations, and to require the parser to search this smaller

space of possible derivations. But none of the proposed methods result in parsing systems which proceed incrementally through the string.⁴

Karttunen (1989) and others have proposed chart-based parsers which directly address the derivational ambiguity problem. For the present purpose, the principal feature of chart parsing — the factoring out of constituents from analyses — turns out to create an encumbrance: The interpreter cannot compare constituents, or arcs, for the purposes of ambiguity resolution. It must compare analyses of the entire prefix so far, which are awkward to compute from the developing chart.

Pareschi and Steedman (1987) propose the following strategy: (which can be taken out of the chart-parsing context of their paper) construct only maximally left-branching derivations, but allow a limited form of backtracking when a locally non-left-branching derivation turns out to have been necessary. For example, when parsing (5), Pareschi and Steedman’s algorithm constructs the left branching analysis for ‘John loves Mary’. When it encounters ‘madly’, it applies >0 in reverse to *solve* for the hidden VP constituent ‘loves Mary’ by *subtracting* the s/vp category ‘John’ from the s category ‘John loves Mary’:

$$(6) \quad \frac{\frac{\frac{\frac{\text{John}}{s/vp} \quad \frac{\text{loves}}{vp/np}}{s/np} \quad \frac{\text{Mary}}{np}}{s} \quad \frac{\text{madly}}{vp \backslash vp}}{\text{vp}} \begin{matrix} >1 \\ >0 \\ <0 \\ >0 \end{matrix}$$

The idea with this ‘revealing’ operation is to exploit the fact that the rules $>n$ and $<n$, when viewed as three-place relations, are functional in all three arguments. That is, knowledge any two of {left constituent, right constituent, result}, uniquely determines the third. There are many problems with the completeness and soundness Pareschi and Steedman’s proposal (Hepple 1987; Niv 1993a). For example, in (7), the category $b \backslash c$ cannot be revealed after it had participated in two combinations of mixed direction: <0 and >0 .

³Discussion of the consequences of this move on the processing of picture noun extractions and ambiguity-related filled-gap effects is omitted for lack of space. See Niv (1993a).

⁴In the case of Hepple’s (1991) proposal, a left-branching normal form is indeed computed. But its computation must be delayed for some words, so it does not provide the interpreter with timely information about the incoming string.

$$(7) \quad \frac{\frac{\frac{a/b \quad \frac{\frac{c \quad d \backslash c}{d} < 0 \quad b \backslash d}{b} < 0 \quad b \backslash c \backslash (b \backslash c)}{b} > 0}{a} \text{ stuck}$$

6 A Proposal

Pareschi and Steedman's idea of lazy parsing is very attractive in the present setting. I propose to replace their unification-based revealing operation with a normal-form based manipulation of the derivation history. The idea is to construct and maintain the maximally incremental, left-branching derivations. (see section 4.) When a constituent such as the VP 'loves Mary' in (5) may be necessary, e.g. whenever the right-most constituent in an analysis is of the form $X \backslash Y$, the next-to-right-most derivation is rewritten to its equivalent right-branching derivation by repeated application the local transformations \rightarrow defined in (8) and (9). The right frontier of the rewritten derivation now provides all the grammatically possible attachment sites.

(8)

$$\frac{\frac{W/X \quad X | Y_1 \cdots | Y_{m-1} / Y_m \quad Y_m | Z_1 \cdots | Z_n}{W | Y_1 \cdots | Y_{m-1} / Y_m} >_m}{W | Y_1 \cdots | Y_{m-1} | Z_1 \cdots | Z_n} >_n \rightarrow \frac{W/X \quad X | Y_1 \cdots | Y_{m-1} / Y_m \quad Y_m | Z_1 \cdots | Z_n}{X | Y_1 \cdots | Y_{m-1} | Z_1 \cdots | Z_n} >_n}{W | Y_1 \cdots | Y_{m-1} | Z_1 \cdots | Z_n} >_{m+n-1}$$

(9)

$$\frac{\frac{Y_m | Z_1 \cdots | Z_n \quad X | Y_1 \cdots | Y_{m-1} \backslash Y_m \quad W \backslash X}{X | Y_1 \cdots | Y_{m-1} | Z_1 \cdots | Z_n} <_n}{W | Y_1 \cdots | Y_{m-1} | Z_1 \cdots | Z_n} <_{m+n-1} \rightarrow \frac{Y_m | Z_1 \cdots | Z_n \quad X | Y_1 \cdots | Y_{m-1} \backslash Y_m \quad W \backslash X}{W | Y_1 \cdots | Y_{m-1} \backslash Y_m} <_m}{W | Y_1 \cdots | Y_{m-1} | Z_1 \cdots | Z_n} <_n$$

Results from the study of rewrite systems (see Klop (1992) for an overview) help determine the computational complexity of this operation:

6.1 A Rewrite System for Derivations

If x is a node in a binary tree let $\lambda(x)$ (resp. $\rho(x)$) refer to its left (right) child.

Any subtree of a derivation which matches the left-hand-side of either (8) or (9) is called a *redex*. The result of replacing a redex by the corresponding right-hand-side of a rule is called the *contractum*. A derivation is in *normal form (NF)* if it contains no redexes. In the following I use the symbol \rightarrow to also stand for the relation over pairs of derivations such that the second is derived from the first by one application of \rightarrow . Let \leftarrow be the converse of \rightarrow . Let \longleftrightarrow be $\rightarrow \cup \leftarrow$. Let \rightarrow^* be the reflexive transitive closure of \rightarrow and similarly, \leftarrow^* the reflexive transitive closure of \leftarrow , and \longleftrightarrow^* the reflexive transitive closure of \longleftrightarrow . Note that \longleftrightarrow^* is an equivalence relation.

A rewrite system is *strongly normalizing (SN)* iff every sequence of applications of \rightarrow is finite.

Theorem 1 \rightarrow is SN⁵

proof Every derivation with n internal nodes is assigned a positive integer score. An application of \rightarrow is guaranteed to yield a derivation with a lower score. This is done by defining functions $\#$ and σ for each node of the derivation as follows:

$$\#(x) = \begin{cases} 0 & \text{if } x \text{ is a leaf node} \\ 1 + \#(\lambda(x)) + \#(\rho(x)) & \text{otherwise} \end{cases}$$

$$\sigma(x) = \begin{cases} 0 & \text{if } x \text{ is a leaf node} \\ \sigma(\lambda(x)) + \sigma(\rho(x)) + \#(\lambda(x)) & \text{otherwise} \end{cases}$$

Each application of \rightarrow decreases σ , the score of the derivation. This follows from the monotonic dependency of the score of the root of the derivation upon the scores of each sub-derivation, and from the fact that locally, the score of a redex decreases when \rightarrow is applied: In figure 2, a derivation is depicted schematically with a redex whose sub-constituents are named a , b , and c . Applying \rightarrow reduces $\sigma(e)$, hence the score of the whole derivation.

in redex:

$$\begin{aligned} \#(d) &= \#(a) + \#(b) + 1 \\ \sigma(d) &= \sigma(a) + \sigma(b) + \#(a) \\ \sigma(e) &= \sigma(d) + \sigma(c) + \#(d) \\ &= \sigma(a) + \sigma(b) + \sigma(c) + \#(b) + 2 \cdot \#(a) + 1 \end{aligned}$$

⁵Hepple and Morrill (1989) Proved SN for a slight variant of \rightarrow . The present proof provides a tighter score function, see lemma 1 below.




Figure 2: Schema for one redex in DRS

in contractum:

$$\begin{aligned}
\sigma(f) &= \sigma(b) + \sigma(c) + \#(b) \\
\sigma(e') &= \sigma(a) + \sigma(f) + \#(a) \\
&= \sigma(a) + \sigma(b) + \sigma(c) + \#(b) + \#(a) \\
&< \sigma(a) + \sigma(b) + \sigma(c) + \#(b) + 2 \cdot \#(a) + 1
\end{aligned}$$

□

Observe that $\#(x)$ is the number of internal nodes in x .

Lemma 1 Given a derivation x , let $n = \#x$. Every sequence of applications of \longrightarrow is of length at most $n(n-1)/2$.⁶

proof By induction on n :

Base case: $n = 1$; 0 applications are necessary.

Induction: Suppose true for all derivations of fewer than n internal nodes. Let $m = \#\lambda(x)$. So $0 \leq m \leq n-1$ and $\#\rho(x) = n-m-1$.

$$\begin{aligned}
\sigma(x) - n(n-1)/2 &= \\
&= \sigma(\lambda(x)) + \sigma(\rho(x)) + \#(\lambda(x)) - n(n-1)/2 \\
&\leq \frac{m(m-1)}{2} + \frac{(n-m-1)(n-m-2)}{2} + m - \frac{n(n-1)}{2} \\
&= (m+1)(m-(n-1)) \\
&\leq 0 \quad \text{recalling that } 0 \leq m \leq n-1
\end{aligned}$$

□

So far I have shown that every sequence of applications of \longrightarrow is not very long: at most quadratic in the size of the derivation. I now show that when there is a choice of redex, it makes no difference which redex one picks. That is, all redex selection strategies result in the same normal form.


A rewrite system is *Church-Rosser (CR)* just in case

$$\forall x, y. (x \longleftarrow y \supset \exists z. (x \longrightarrow z \wedge y \longrightarrow z))$$

A rewrite system is *Weakly Church-Rosser (WCR)* just in case

$$\forall x, y, w. (w \longrightarrow x \wedge w \longrightarrow y) \supset \exists z. (x \longrightarrow z \wedge y \longrightarrow z)$$

⁶Niv (1994) shows by example that this bound is tight.



Arrows are annotated by the substructure to which they are applied

Figure 3: Why \longrightarrow is weakly Church-Rosser

Lemma 2 \longrightarrow is WCR.

proof Let w be a derivation with two distinct redexes x and y , yielding the two distinct derivations w' and w'' respectively. There are a few possibilities:

case 1: x and y share no internal nodes. There are three subcases: x dominates y (includes y as a subconstituent), x is dominated by y , or x and y are incomparable with respect to dominance. Either way, it is clear that the order of application of \longrightarrow makes no difference.

case 2: x and y share some internal node. Without loss of generality, y does not dominate x . There exists a derivation z such that $w' \longrightarrow z \wedge w'' \longrightarrow z$. This is depicted in figure 3. (Note that all three internal nodes in figure 3 are of the same rule direction, either $>$ or $<$.)

□

Lemma 3 (Newman) $\text{WCR} \wedge \text{SN} \supset \text{CR}$.

Theorem 2 \longrightarrow is CR.

proof From theorem 1 and lemmas 2 and 3. □

Therefore any maximal sequence of applications of \longrightarrow will lead to the normal form⁷. We are free to select the most efficient redex selection scheme. From lemma 1 the worst case is quadratic. Niv (1994) shows that the optimal strategy, of ap-

⁷Assuming, as is the case with extant CCG accounts, that constraints on the applicability of the combinatory rules do not present significant roadblocks to the derivation rewrite process.

plying \rightarrow closest as possible to the root, yields \rightarrow applications sequences of at most n steps.

Note that all that was said in this section generalizes beyond CCG derivations to any associative algebra.

6.2 Discussion

Given the rightmost subconstituent recovered using the normal form technique above, how should parsing proceed? Obviously, if the leftward looking category which precipitated the normal form computation is a modifier, i.e. of the form $X \backslash X$, then it ought to be combined with the recovered constituent in a form analogous to Chomsky adjunction. But what if this category is not of the form $X \backslash X$? For example, should the parser compute the reanalysis in (10)?

$$\begin{array}{rcl}
 (10) & \frac{\frac{a/b \quad b/c}{a/c} >1 \quad c/d \quad s \backslash (a/b) \backslash (b/d)}{a/d} >1 \\
 & \frac{a/b \quad \frac{b/c \quad c/d}{b/d} >1 \quad s \backslash (a/b) \backslash (b/d)}{s \backslash (a/b)} <0 \\
 & s <0
 \end{array}$$

Ascribing the same non-garden-path status to the reanalysis in (10) that we do to (6) would constitute a very odd move: Before reanalysis, the derivation encoded the commitment that the $/b$ of the first category is satisfied by the b of the b/c in the second category. This commitment is undone in the reanalysis. This is an undesirable property to have in a computational model of parsing commitment, as it renders certain revisions of commitments easier than others, without any empirical justification. Furthermore, given the possibility that the parser change its mind about what serves as argument to what, the interpreter must be able to cope with such non-monotonic updates to its view of the analysis so far — this would surely complicate the design of the interpreter.⁸ Therefore, constituents on the right-frontier of a right-normal-form should only combine with ‘endocentric’ categories to their right. The precise definition of ‘endocentric’ depends on the semantic formalism used — it certainly includes post-head modifiers, and might also include coordination.

⁸I am indebted to Henry Thompson for a discussion of monotonicity.

Stipulating that certain reanalyses are impossible immediately makes the parser ‘incomplete’ in the sense that it cannot find the analysis in (10). From the current perspective of identifying garden paths, this incompleteness is a desirable, even a necessary property. In (10), committing to the composition of a/b and b/c is tantamount to being led down the garden path. In a different sense, the current parser is complete: it finds all analyses if the Viable Analysis Criterion and the interpreter never discard any analyses.

7 Conclusion

The current proposal shifts some of the burden traditionally associated with the parser to other components of the human cognitive faculty: the interpreter resolves ambiguity, and an acquired skill removes ‘garbage’ analyses from the parser’s memory — solving the so-called spurious ambiguity problem, as well as effectively applying grammar-global constraints traditionally computed by top-down techniques or grammar compilation. The resultant parser adheres to the desideratum that it be a generic search algorithm for the grammar formalism, provided the definition of CCG explicitly includes the notion of ‘derivation’ and explicates the truth-conditional equivalence relation. Such inclusions have indeed been proposed (Steedman 1990).

Bibliography

- Aho, Alfred and S. C. Johnson. 1974. LR Parsing. *ACM Computing Surveys* 6(2):99–124.
- Altmann, Gerry T., Alan Garnham, and Judith A. Henstra. 1994. Effects of syntax in human sentence parsing: Evidence against a structure-based proposal mechanism. *Journal of Experimental Psychology: Learning, Memory and Cognition* 20(1):1–8.
- Dowty, David. 1988. Type Raising, Functional Composition, and Non-Constituent Conjunction. In Richard T. Oehrle, Emmon Bach, and Deirdre Wheeler (Eds.), *Categorial Grammars and Natural Language Structures*. Reidel.
- Earley, Jay. 1970. An Efficient Context-Free Parsing Algorithm. *Communications of the Association for Computing Machinery* 13:94–102.
- Hepple, Mark R. 1987. Methods for Parsing Combinatory Grammars and the Spurious Ambiguity Problem. Master’s thesis, University of Edinburgh.

- Hepple, Mark R. 1991. Efficient Incremental Processing with Categorical Grammar. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, 79–86.
- Hepple, Mark R. and Glyn V. Morrill. 1989. Parsing and Derivational Equivalence. In *Proceedings of the Annual Meeting of the European Chapter of the Association for Computational Linguistics*.
- Karttunen, Lauri. 1989. Radical Lexicalism. In Mark Baltin and Anthony S. Kroch (Eds.), *Alternative Conceptions of Phrase Structure*, 43–65. Chicago: University of Chicago Press.
- Kawamoto, Alan and William Farrar. 1993. The Return of Visiting Relatives: Pragmatic Effects in Sentence Processing. *Quarterly Journal of Experimental Psychology* 46A(3):463 – 487.
- Klop, Jan W. 1992. Term Rewrite Systems. In Samson Abramsky, Dov M. Gabbay, and T. S. E. Maibaum (Eds.), *Handbook of Logic in Computer Science*, Vol. 2, 1 – 116. Oxford: Clarendon Press.
- König, Esther. 1989. Parsing as Natural Deduction. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, 272–279, June.
- Lambek, Joachim. 1958. The Mathematics of Sentence Structure. *American Mathematical Monthly* 65:154–169.
- MacDonland, Maryellen, Adam Just, and Patricia Carpenter. 1992. Working Memory Constraints on the Processing of Syntactic Ambiguity. *Cognitive Psychology* 24:56–98.
- Nicol, Janet L. and Martin J. Pickering. 1993. Processing Syntactically Ambiguous Sentences: Evidence from Semantic Priming. *Journal of Psycholinguistic Research*.
- Niv, Michael. 1993a. *A Computational Model of Syntactic Processing: Ambiguity Resolution from Interpretation*. PhD thesis, University of Pennsylvania. (electronically: cmp-lg/9406029 or ftp://ftp.cis.upenn.edu/pub/-ircs/tr/93-27.ps).
- Niv, Michael. 1993b. Resolution of Syntactic Ambiguity: the Case of New Subjects. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, Hillsdale, NJ: Lawrence Erlbaum Associates. (electronically: cmp-lg/9406028 or ftp://ftp.cis.upenn.edu/pub/niv/cogsci93.ps).
- Niv, Michael. 1994. The complexity of normal form rewrite sequences for Associativity. Technical Report LCL 94-6, Computer Science Department, Technion. (electronically: cmp-lg/9406030 or ftp://ftp.cis.upenn.edu/pub/niv/rewrite.ps).
- Pareschi, Remo and Mark J. Steedman. 1987. A Lazy Way to Chart Parse with Combinatory Grammars. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*.
- Shieber, Stuart M. and Mark Johnson. 1993. Variations on Incremental Interpretation. *Journal of Psycholinguistic Research*. (to appear).
- Steedman, Mark J. 1987. Combinatory Grammars and Parasitic Gaps. *Natural Language and Linguistic Theory* 5:403–439.
- Steedman, Mark J. 1990. Gapping as Constituent Coordination. *Linguistics and Philosophy* 13:207–264.
- Steedman, Mark J. 1994. Grammars and Processors. In Hans Kamp and Christian Rohrer (Eds.), *Aspects of Computational Linguistics*. Springer Verlag. (to appear).
- Trueswell, John C. and Michael K. Tanenhaus. 1991. Tense, Temporal Context and Syntactic Ambiguity Resolution. *Language and Cognitive Processes* 6:303–338.
- Trueswell, John C. and Michael K. Tanenhaus. 1994. Toward a Lexicalist Framework of Constraint-Based Syntactic Ambiguity Resolution. In Charles Clifton, Lyn Frazier, and Keith Rayner (Eds.), *Perspectives on Sentence Processing*. Hillsdale, NJ: Lawrence Erlbaum Associates.